



Cover Letter for the 2024 PESO Project Report

Partnering for Scientific-Software Ecosystem Opportunities (PESO)

Dear Reader,

We are pleased to share with you the 2024 PESO Project Report. Science increasingly relies on powerful computers and the software that runs on them. From modeling the basic science of nanomaterials and simulating earthquakes to advancing medical research and energy technologies, this software plays a critical role in solving society's biggest challenges. The PESO Project exists to support scientific discovery through better software.

PESO is dedicated to stewarding and advancing the current and future ecosystem of HPC-AI software and works to make this software more reliable, easier to use, and accessible to scientists and engineers across the country. Our team includes experts from national laboratories, universities, and private industry who collaborate to improve the quality and sustainability of scientific computing tools.

In 2024, our project focused on:

- Improving tools that help scientists install and manage software more easily, especially in complex computing environments like supercomputers and cloud platforms, including the Department of Energy's flagship Exascale systems capable of more than a billion-billion operations per second.
- Expanding and refining the E4S software stack, which is a curated collection of scientific software used by researchers in many fields.
- Enhancing user and developer experiences, making it easier for scientists and engineers to use these tools effectively.
- Measuring the impact of scientific software, helping to ensure that public investments support high-quality and high-value research tools.
- Fostering a strong scientific software community, including the support of early-career scientists and career paths in software development for science.

Through these efforts, PESO contributes to the U.S. Department of Energy's mission to advance scientific discovery, while also helping build a modern, flexible, and comprehensive computing infrastructure that benefits a broad spectrum of users.

We invite you to explore the full report. We hope it gives you a sense of the important work being done to ensure that scientific software continues to serve as a powerful engine for discovery, innovation, and progress.

Sincerely,
The PESO Project Team



2024 PESO Project Report

PESO: Partnering for Scientific-Software Ecosystem Opportunities

December 17, 2024

Michael Heroux	ParaTools, Inc., PI
Lois Curfman McInnes	Argonne National Laboratory, PI
Todd Gamblin	Lawrence Livermore National Laboratory
Sameer Shende	University of Oregon
James Willenbring	Sandia National Laboratories
Satish Balay	Argonne National Laboratory
Roscoe Bartlett	Sandia National Laboratories
Keith Beattie	Lawrence Berkeley National Laboratory
Greg Becker	Lawrence Livermore National Laboratory
David Bernholdt	Oak Ridge National Laboratory
Sam Browne	Sandia National Laboratories
Hannah Cohoon	Lawrence Berkeley National Laboratory
Tamara Dahlgren	Lawrence Livermore National Laboratory
Berk Geveci	Kitware, Inc.
William Godoy	Oak Ridge National Laboratory
Elsa Gonsiorowski	Lawrence Livermore National Laboratory
Patricia Grubel	Los Alamos National Laboratory
Mahantesh Halappanavar	Pacific Northwest National Laboratory
Bill Hoffman	Kitware, Inc.
Damien Lebrun-Grandie	Oak Ridge National Laboratory
Mary Ann Leung	Sustainable Horizons Institute
Xiaoye Sherry Li	Lawrence Berkeley National Laboratory
Daniel Martin	Lawrence Berkeley National Laboratory
Mark C. Miller	Lawrence Livermore National Laboratory
Marco Minutoli	Pacific Northwest National Laboratory
Patrick O'Leary	Kitware, Inc.
Erik Palmer	Lawrence Berkeley National Laboratory
Suzanne Parete-Koon	Oak Ridge National Laboratory
Lavanya Ramakrishnan	Lawrence Berkeley National Laboratory
Rajeev Thakur	Argonne National Laboratory
Mikhail Titov	Brookhaven National Laboratory
Matteo Turilli	Brookhaven National Laboratory
Terece Turton	Los Alamos National Laboratory
Ulrike Meier Yang	Lawrence Livermore National Laboratory
Hui Zhou	Argonne National Laboratory

Table of Contents

<i>Executive Summary</i>	4
Key Highlights during 2024	4
PESO and CASS Collaboration.....	4
E4S Enhancements and Adoption:	4
Spack Enhancements and Adoption:	5
Software Ecosystem Integration:	5
User and Developer Experience (UDX):.....	5
Impact Frameworks:	5
Community Engagement:.....	5
1. Introduction	7
2. PESO and CASS: A Multifaceted Collaboration	8
3. E4S Enhancements and Adoption Report	9
3.1. Scope	9
3.2. Team	9
3.3. Achievements	9
3.4. Current Status	10
3.5. Plans	10
4. Spack Enhancements and Adoption	11
4.1 Scope	11
4.2 Team	11
4.3 Achievements	11
4.4 Current Status	12
4.5 Plans	12
5. Software Ecosystem Integration	13
5.1 Scope	13
5.2 Team	13
5.3 Achievements	14
5.4 Current Status	15
5.5 Plans	15
6. User and Developer Experience (UDX)	16
6.1 Scope	16
6.2 Team	16

2024 PESO Project Report

6.3 Achievements	17
6.4 Current Status.....	17
6.5 Plans.....	18
<i>7. Impact Framework.....</i>	<i>18</i>
7.1 Scope	18
7.2 Team.....	18
7.3 Achievements	19
7.4 Current Status.....	20
7.5 Plans.....	20
<i>8. PESO Community Engagement.....</i>	<i>21</i>
8.1 Scope	21
8.2 Team.....	22
8.3 Achievements	22
8.4 Current Status.....	23
8.5 Plans.....	23
<i>9. Conclusion.....</i>	<i>23</i>

2024 PESO Project Report

Executive Summary

Computing is a foundational capability in almost all areas of scientific discovery. For problems that demand high-end computing resources, stewarding and advancing the software ecosystem established by the US Exascale Computing Project (ECP) is a fundamental responsibility, building on ECP investments.

The [PESO project](#), in collaboration with the [Consortium for the Advancement of Scientific Software \(CASS\)](#), is building a software ecosystem that enables foundational scientific research and advancements. PESO is advancing scientific software development, adoption, and community engagement across multiple thrust areas. These areas include Spack enhancements, E4S (Extreme-scale Scientific Software Stack) improvements, community engagement, software ecosystem integration, user and developer experience (UDX), impact evaluation frameworks, and community development.

PESO's efforts align with the DOE mission of advancing fundamental scientific discovery by improving the quality, sustainability, and adoption of scientific software, making it accessible and usable for a diverse community of users in both high-performance computing (HPC) and artificial intelligence (AI) domains. PESO emphasizes partnerships, continuous integration, community support, and outreach activities that ensure sustained impact and growth in scientific software quality, career pathways, and infrastructure stability—all to foster a software ecosystem that enables foundational scientific research and advancements.

Key Highlights during 2024

PESO and CASS Collaboration

- Played a key role in establishing CASS; Lois Curfman McInnes and David Bernholdt were facilitators; senior PESO members drafted the charter and bylaws, adopting Martha's Rules for consensus-based decision-making.
- Led the CASS Integration, Impact Framework, User-Developer Experience, and Next-generation Scientists Working Groups.
- Organized the first CASS 'Birds of a Feather' (BoF) days.
- Hosted a PESO annual planning meeting at Argonne National Laboratory, along with a CASS annual planning meeting, emphasizing crosscutting strategic planning.
- Sponsored CASS leaders David Bernholdt (CASS Chair) and Terry Turton (Vice Chair).

E4S Enhancements and Adoption:

- Expanded the number of supported software packages, including over 140 HPC and AI tools.
- Expanded the AI packages included in E4S with support for Large Language Models (LLMs) and NVIDIA NeMo™ for GPUs.

2024 PESO Project Report

- Improved Spack-based deployment methods and enhanced documentation and user guides.
- Hosted workshops and webinars, reaching over 500 participants, and supported projects such as Trilinos, Kokkos, Viskores, and ExaWind for continuous integration and testing.
- Collaborated with commercial cloud platforms (AWS, Google Cloud, and Azure) to expand adoption, resulting in over 1 million downloads of E4S containers.
- Developed a draft strategy for a cohesive AI and ModSim software ecosystem, submitted to the RFI for FASST.

Spack Enhancements and Adoption:

- Participated in efforts for major releases (v0.21.0 and v0.22.0) with enhancements to the concretizer, testing interface, and dependency management system.
- Participated in an efforts-migration plan for the upcoming "compilers as dependencies" feature, ensuring backward compatibility.
- Established a new Continuous Integration (CI) working group within the High-Performance Software Foundation (HPSF).

Software Ecosystem Integration:

- Focused on improving quality, sustainability, and interoperability across the scientific software ecosystem.
- Developed a minimal validating test (Spack Smoke Test) for E4S products and optimized CI testing across ecosystem levels.

User and Developer Experience (UDX):

- Established a UDX working group to improve the usability, accessibility, and productivity of PESO's tools and technologies.
- Raised awareness of user experience in the PESO and CASS communities, sharing resources on user experience methods and value.
- Planned a UDX research study to investigate challenges and identify potential solutions, aiming to gather best practices and communicate their impact.

Impact Frameworks:

- Established the Impact Framework Working Group to develop adaptable metrics for evaluating the impact of scientific software.
- Designated as a CASS-level working group, focusing on productivity, sustainability, and quality metrics to inform funding decisions.
- Initiated pilot testing of metrics to assess software impact and align efforts with broader DOE evaluation initiatives.

Community Engagement:

- Organized and delivered the CASS Community BOF Days.

2024 PESO Project Report

- Coordinated the Better Scientific Software (BSSw) Fellowship Program (jointly sponsored by ASCR, NNSA, and NSF), selecting six Fellows and six Honorable Mentions for the 2024 cohort and five Fellows and five Honorable Mentions for the 2025 cohort.
- Secured partial funding for the BSSw Fellows Program for the next several years while working toward sustainable sponsorship.
- Published more than 40 articles on the BSSw.io website to disseminate best practices in scientific software development.

2024 PESO Project Report

Abstract

Computing is a foundational capability in almost all areas of scientific discovery. The PESO project is dedicated to enhancing scientific software through a variety of collaborative initiatives that span multiple areas, including package management tools like Spack, the E4S software stack, scientist training, user and developer experience, and frameworks to evaluate the impact of funded software projects. The project partners with organizations such as the National Science Foundation (NSF), the National Nuclear Security Administration (NNSA), and academic collaborators to build, enhance, and sustain an ecosystem that fosters scientific innovation and computational research. This report provides an overview of the project's efforts, achievements, and plans, illustrating PESO's role in improving software infrastructure and advancing the high-performance computing community.

1. Introduction

Computing is a foundational capability in almost all areas of scientific discovery.¹ For problems that demand high-end computing resources, stewarding and advancing the software ecosystem established by the US Exascale Computing Project (ECP) is a fundamental responsibility, building on ECP investments.

The PESO project is a comprehensive initiative to foster scientific software development, integration, and sustainability across a broad spectrum of high-performance computing (HPC) applications. As a key contributor to the Consortium for the Advancement of Scientific Software (CASS), PESO plays a central role in various community-driven projects that support the creation of impactful scientific software tools and methodologies. The relationship between PESO and CASS is built upon shared goals, a collaborative approach to decision-making, and a commitment to supporting the scientific software community.

The PESO project's scope includes creating environments that foster reproducibility, accessibility, and usability of scientific software, making it possible for researchers, engineers, and developers to conduct and advance their work efficiently. PESO's initiatives, such as Spack and E4S enhancements, support efficient software deployment, while community engagement thrusts ensure that a diverse group of stakeholders—including vendors, researchers, and various institutions—actively participate in the project's evolution. Moreover, PESO's focus on software ecosystem integration, as well as user and developer experience, aims to reduce barriers to adoption and improve the productivity of scientific computing tools.

¹ *The Science of Scientific Software Development and Use* (<https://doi.org/10.2172/1846009>); *AI for Science, Energy, and Security* (<https://doi.org/10.2172/1986455>)

2024 PESO Project Report

2. PESO and CASS: A Multifaceted Collaboration

The PESO project has cultivated a multifaceted relationship with the Consortium for the Advancement of Scientific Software (CASS). One of the fundamental connections between these organizations lies in the composition of PESO itself. Many PESO team members are also part of other Software Stewardship Organizations (SSOs), and these SSOs, along with PESO, are the founding members of CASS.

PESO has significantly contributed to the formation of CASS, particularly in terms of staffing and leadership. Two key PESO team members, Lois Curfman McInnes, and David Bernholdt, were instrumental in guiding the formation of CASS as they served as two of the three facilitators of this process. Other senior members of PESO also contributed significantly to drafting the CASS charter and bylaws. Of note is the adoption of Martha's Rules, which encourage consensus-based decision-making. This commitment to collaboration and engagement is also reflected in the expectation that CASS meetings are open to all members of the broader CASS community, even if steering committee membership is limited to one representative per SSO.

PESO team members also provide significant leadership to several CASS working groups:

- **Integration Working Group:** Focuses on coordination across teams to facilitate the building, integrating, and testing of complementary software products for improved user experience when working with our software as a portfolio.
- **Impact Framework:** Focuses on establishing a flexible framework for understanding and improving software product impact.
- **User-developer experience (UDX):** Focuses on improving awareness and use of techniques for user and developer experience with scientific software, and another focused on the definition and advancement of software ecosystems for more effective and efficient use of libraries and tools in scientific applications.
- **National Competitiveness Working Group:** Focuses on engaging the next generation of scientists with pathways to impactful HPC & AI careers.

PESO's role in fostering community engagement extended to organizing and sponsoring the first CASS community "Birds of a Feather" (BoF) days, held in June 2024. PESO led the coordination efforts, collecting input from other CASS teams and executing the BoF sessions successfully. PESO plans to continue providing this critical coordinating role annually, fostering dialogue and collaboration across the scientific software community. Planning is under way for CASS BOF Days, to be held in early 2026.

Additionally, PESO organized its annual planning meeting at Argonne National Laboratory in July 2024 and extended an invitation to members of the other SSOs for a dedicated CASS session held the day after the PESO meeting. This kind of coordinated, collocated meeting enabled PESO to conduct its own business while creating opportunities for alignment and coordination with CASS and the other SSOs. Planning is under way for similar sessions in July 2025.

2024 PESO Project Report

Finally, PESO is a significant supporter of CASS leadership. PESO sponsors the time of David Bernholdt, who serves as the CASS Chair, and Terry Turton, Vice Chair. Since CASS does not have direct funding for its operation, these leadership roles rely on external funding, and PESO's support has been essential for ensuring CASS's continued growth and effectiveness.

3. E4S Enhancements and Adoption Report

3.1. Scope

PESO's thrust on E4S (Extreme-scale Scientific Software Stack, <https://e4s.io>) Enhancements and Adoption aims to expand the capabilities, usability, and adoption of the E4S ecosystem. This work focuses on providing a curated software stack to facilitate advances in high-performance scientific applications, addressing the diverse needs of the scientific community by offering a curated collection of build-from-scratch, pre-built, and optimized software libraries and tools. This work integrates into the broader ecosystem of scientific computing by fostering collaboration and simplifying access to tools essential for research and development at scale. E4S is a key part of PESO's efforts because it is key to improving accessibility, productivity, and sustainability of software infrastructure, which underpin scientific discoveries.

3.2. Team

- Overall Lead: Michael Heroux
- Technical Lead: Sameer Shende
- Integration Lead: James Willenbring
- Integration Team:
 - Luke Peyralans
 - Eugene Walker
 - Wyatt Spear
- Key community contributors: Experts from computing facilities at Argonne, Oak Ridge, and Lawrence Berkeley National Laboratories, as well as external collaborators from academic institutions and industry partners
- Other contributors: Many people who have provided bug fixes, documentation improvements, and porting support

3.3. Achievements

- Expanded the number of supported software packages, including over 140 HPC and AI tools.
- Completed enhancements to Spack-based deployment methods, enabling easier installation on different architectures, including cloud-based environments.
- Improved documentation and user guides make it easier for new users to start with E4S.
- Hosted several hands-on workshops and webinars, reaching over 500 participants, to promote the capabilities of E4S and train new users.

2024 PESO Project Report

- Worked with the San Diego Supercomputer Center’s Earthquake Simulation team who used E4S to simplify dependency management and enhance performance analysis for the AWP-ODC application, improving load balancing and workflow complexity.
- Worked with the NOAA's Cloud Workflow to leverage E4S to create Amazon Machine Images (AMIs) and support containerized workflows, resulting in over 1 million downloads of E4S containers.
- Supported continuous integration effort for projects such as Trilinos and ExaWind, ensuring software quality across nine different GPU architectures enhancing developer confidence in compatibility.
- Delivered E4S images on AWS and Google Cloud, with partnerships established for integrating E4S into commercial AI and HPC offerings, expanding adoption.
- Executed nearly 8 million jobs using E4S across various architectures, demonstrating its versatility and robust performance.
- Established regular CI of Kokkos with GitLab pipelines using E4S containers providing access to advanced GPU architectures and enhanced software reliability.

3.4. Current Status

- Collaboration with HPC centers: Gathering feedback from HPC centers to prioritize user-requested features.
- Improving container support: Making deployment in cloud-native settings more seamless.
- Expanding CI/CD capabilities: Ensuring software quality across multiple architectures, including AMD, Intel, and NVIDIA GPUs and conducting testing on nine different GPU architectures to maintain software quality and compatibility.
- Challenges due to reduced funding: Addressing challenges in maintenance and release frequency following the conclusion of the Exascale Computing Project (ECP). Limited resources have impacted updates on major systems, but the team remains committed to mitigating these issues.
- Community support and collaboration: Emphasizing the importance of sustained resources to maintain software quality and security through community-driven efforts.

3.5. Plans

In FY25, the team will prioritize further expanding the range of software packages supported by E4S, focusing on enhancing support for AI and machine learning workflows. Another central area of focus will be integration with workflow management systems to provide end-to-end solutions for scientific pipelines. By 2028, the goal is for E4S to be the default software stack used in at least 75% of DOE facilities, offering comprehensive support for various scientific applications.

The long-term vision includes making E4S a go-to solution for traditional HPC and AI-driven workloads, bridging the gap between these computing paradigms. The PESO team

2024 PESO Project Report

produced a draft strategy for a cohesive software stack for AI and Modeling & Simulation (AI/ModSim), submitted as a response to the FASST RFI.²

4. Spack Enhancements and Adoption

4.1 Scope

The focus of this thrust area is to enhance and promote the adoption of Spack (<https://spack.io>), an open-source package management tool tailored for high-performance computing (HPC). The primary aim is to extend its capabilities to better serve the HPC and broader scientific software ecosystem by improving scalability, ease of use, and integration with other tools. This effort fits into the wider ecosystem as Spack provides critical infrastructure for managing complex software environments, making it a key tool in modernizing HPC workflows and ensuring reproducibility. Spack is a key part of PESO's efforts to create a cohesive, user-friendly, and sustainable software environment that accelerates scientific discovery across projects and institutions.

4.2 Team

The Spack enhancements and adoption efforts are led by a collaborative team involving both PESO members and external contributors. PESO staff involved in Spack's efforts are:

- Todd Gamblin: Spack PI
- Tamara Dahlgren: Senior Spack developer
- Greg Becker: Senior Spack developer

4.3 Achievements

In FY24, the Spack team successfully delivered two major releases: v0.21.0 and v0.22.0. Version 0.21.0 introduced enhancements to the concretizer, a new standalone testing interface, and a more stable dependency management system. Version 0.22.0 expanded Spack's capabilities for handling compilers, improved external package detection, and added support for libc-based compatibility on Linux systems. Additionally, several bug fix releases—v0.21.1, v0.21.2, v0.21.3, v0.22.1, and v0.22.2—addressed issues and maintain the stability of the project.

Community support has been a significant focus this year. The team overhauled the testing support in Spack packages, porting many packages using legacy testing logic to the new standalone test logic. Spack has also continued its strong community engagement by merging hundreds of pull requests each month to incorporate the latest software versions and user contributions. These efforts ensure that Spack remains a robust and up-to-date

² *Toward a Cohesive AI and Simulation Software Ecosystem for Scientific Innovation:*
<https://arxiv.org/abs/2411.09507>

2024 PESO Project Report

tool that serves the needs of its broad user base.

Another significant activity has been developing and implementing a migration plan for the upcoming "compilers as dependencies" feature. This has required migrating user configuration files to accommodate the new functionality. Additionally, the team has begun overhauling the internal representation of dependency graphs. This ongoing work involves complex changes that require careful backward compatibility considerations to ensure a smooth transition. To help users adapt, the team released two versions of Spack in November: one for version 0.23.0 and a release candidate for the new compilers feature, which will be fully integrated by June next year.

4.4 Current Status

After finalizing Spack 0.23.0 and a release candidate featuring the new "compilers as dependencies" support, we are continuing to actively test to ensure continuing compatibility across a wide range of platforms and user configurations. The overhaul of Spack's dependency graph representation is in progress, and the team is working to ensure that this complex refactoring is fully backward compatible, with detailed testing and documentation to guide users through the transition.

4.5 Plans

In FY25, the primary focus will be on responding to feedback from the compiler dependencies release candidate and incorporating any necessary adjustments before a full integration in mid-year. The team will prioritize gathering user input, testing the new feature across diverse environments, and ensuring a smooth adoption process for all Spack users. This work will involve refining migration tools, addressing potential compatibility issues, and enhancing the user documentation to support this significant change.

In parallel, the team will collaborate closely with the Spack Technical Steering Committee (TSC) to identify and refine project-wide priorities. Regular meetings and discussions will help align the team's development efforts with the broader goals of the Spack community. This iterative process will ensure that Spack continues to evolve in a way that meets both DOE user needs and long-term strategic objectives.

Another key priority will be a significant overhaul of Spack's build cache infrastructure. The goal is to make build caches more portable, enabling most Spack installations to fetch most dependencies as binaries rather than building from source. The team aims to enable build caches by default, making Spack installations faster and more reliable for a broader range of users. This work will require modifications to the build cache system for content-based addressing to enable build caches to leverage each other more easily.

Lastly, the Spack project has just started a new Continuous Integration (CI) working group within the High-Performance Software Foundation (HPSF). The goal is to leverage Spack's existing CI system as a foundation for building an HPC-oriented CI service for other HPSF projects. This work involves documenting the tagging scheme for runners on the Frank

2024 PESO Project Report

cluster at the University of Oregon and creating tutorials for projects to connect to this GitLab-based CI system. The ultimate vision is to enable projects to easily integrate with the system via a GitHub app, allowing them to leverage the same powerful CI capabilities developed for Spack. The team is currently onboarding projects such as Viskores and Kokkos onto the service at gitlab.spack.io, with plans to migrate to gitlab.hpsf.io in the coming months, making it easier for other projects to adopt these capabilities.

5. Software Ecosystem Integration

5.1 Scope

The PESO Software Ecosystem Integration Thrust Area focuses on efforts at the levels of scientific software ecosystem and individual product communities, including on-node and inter-node programming systems, tools, data and visualization, math libraries, and workflows. The goals are to improve quality, sustainability, delivery, and interoperability across the ecosystem. The dual focus on ecosystem and product levels is essential due to the challenges of creating a cohesive ecosystem that includes product communities with different needs and varying levels of maturity. The efforts in this thrust area are often synergistic with initiatives from other PESO focus areas, such as the Impact Framework and User-Developer Experience (UDX).

A common objective across Integration Thrust Area initiatives is to move activities to higher levels to improve efficiency. For example, certain testing can be transferred from the product community to the ecosystem level, or redundancy can be eliminated in porting efforts. The overarching goal is to sustain and continue improving the interoperability, delivery, and quality of the software technologies ecosystem under ECP, emphasizing community efforts.

The thrust area covers the following product families:

- On-node programming systems: OpenMP/OpenACC, LLVM, Kokkos, Fortran, HIP, SYCL.
- Inter-node programming systems: MPICH, OpenMPI, GasNet.
- Tools: Darshan, Dyninst, HPCToolkit, PAPI, TAU, and others in collaboration with the Software Tools Ecosystem Project (STEP).
- Data & Visualization: DAV SDK, including ADIOS, Viskores, Catalyst, ParaView, DIY, Ascent, HDF5, PNetCDF, VisIt, ZFP.
- Math Libraries: xSDK, which includes multiple math libraries.
- Workflows: RADICAL-Cybertools, Parsl, Flux, and other workflow technologies.

5.2 Team

- Jim Willenbring: PESO Integration Coordinator
- Damien Lebrun-Grandie: On-node programming systems (in collaboration with S4PST)
- Hui Zhou: Inter-node programming systems (in collaboration with S4PST)
- Bill Hoffman: Tools (in collaboration with STEP)

2024 PESO Project Report

- Satish Balay: Math libraries (in collaboration with FASTMATH)
- Patrick O’Leary: Data & visualization (in collaboration with RAPIDS)
- Matteo Turilli: Workflows (in collaboration with SWAS)
- Sam Browne: NNSA software (funded by NNSA)
- Sameer Shende: Tools STEP leader
- Berk Geveci: Partnership Coordinator for DAV, RAPIDS
- David Bernholdt
- Mike Heroux
- Junchao Zhang
- Phil Carns
- Rajeev Thakur
- Terry Turton
- Todd Gamblin
- William Godoy
- Tammy Dahlgren

5.3 Achievements

- Established common objectives through a better understanding of product community and ecosystem needs.
- Converted current product stand-alone tests to the new Spack testing API, with several also being changed to address issues testing at facilities.
- Progressed planning for continuous integration (CI) across levels (ecosystem, product community, product).
- Developed a minimal validating test (Spack Smoke Test) as required by Policy 2 of the E4S Community Policies to provide essential confidence that a user can compile, install, and run every E4S member package.
- DAV integration team: Transitioned from using a Spack meta-package to utilizing a Spack environment template, improving maintainability by correctly integrating the DAV SDK into Spack.
- Optimized ParaView builds to support Spack CI, which reduced the number of separate builds required.
- Held meetings and discussions with product teams for tools integration, such as Darshan, Dyninst, HPCToolkit, PAPI, and TAU.
- Migrated logic from the current DAV meta-package to Spack package configuration files and Spack environment files, improving ease of use.
- Advanced DAV SDK products, including adding ANARI rendering to VTK, upgrading ADIOS with NumPy, and supporting accelerator-based data parallelism in Viskores.
- Integration of Kokkos SYCL backend and enhancement of Kokkos’ experimental OpenACC backend.
- Started coding Spack packages for RADICAL-Cybertools, Parsl, and Flux for workflow integration.

2024 PESO Project Report

5.4 Current Status

Goals for FY25 feature ongoing activities such as transitioning from using a Spack meta-package to utilizing a Spack environment template with configuration files. Efforts are underway to improve continuous integration (CI) practices, focusing on enhanced testing across levels (ecosystem, product community, product) and reducing redundancy in porting efforts.

Facility-level integration CI testing remains an area of improvement. Facility-provided CI services via Gitlab CI are available in most DOE facilities, but facility integration CI testing needs to be more common and present. Most projects have nightly CI testing and per-merge request CI testing, but they often use custom testing harnesses not immediately transferable to Spack testing.

Integration status for specific product families:

- On-node programming systems: Kokkos' automated testing is being set up with gitlab.spack.io to run on platforms such as the University of Oregon's Frank Cluster.
- Inter-node programming systems: Facility CI testing for MPICH has been explored using facility-provided Gitlab CI services, particularly with the Argonne Leadership Computing Facility (ALCF).
- Tools integration: Meetings will be held with the University of Oregon and product teams to improve CI practices across tools products in STEP.
- Data & visualization: The DAV integration team has improved facility-level CI testing, especially for platforms like Frontier and Aurora.
- Math Libraries: Collaboration continues through xSDK, with integration working groups established to enhance math libraries within the E4S stack.
- Workflows: Integrating workflow technologies like RADICAL-Cybertools, Parsl, and Flux is underway, focusing on integration into E4S.

5.5 Plans

Plans for FY25:

- Ensuring that each participating product has meaningful stand-alone tests for Spack and E4S testing.
- Expanding the set of coordinated CI builds.
- Developing an initial implementation of integrated CI testing on one platform.
- Improving CI-related workflows and associated documentation.
- Increasing interoperability between packages and package options exercised by ecosystem-level CI smoke tests.
- Developing CI practices that can adapt to different goals, including developers, users, and system integrators, while avoiding silos in testing across projects.
- Establishing facility-level CI support for runtime integration, starting with MPICH and extending to OpenMPI and GasNet.

2024 PESO Project Report

- Expanding Kokkos automated testing on various platforms, including Frank's cluster and Leadership Computing Facilities systems.
- Advancing the DAV SDK and integrating products like Catalyst, ParaView, VTK, ADIOS, HDF5, PNetCDF, and VisIt into CI workflows.

By 2028, the vision is to:

- Achieve a regularly practiced system integration CI across E4S projects.
- Establish facility-level CI support and promote an integration CI culture that serves the needs of developers, users, and system integrators.
- Regularly practice system integration CI across E4S projects, establishing facility practice of supporting and administering system integration CI.
- Establish recommended practices with sample projects and workflows.
- Deliver production-grade and long-term maintained Spack packages for workflow technologies like RADICAL-Cybertools, Parsl, and Flux, integrated within E4S.

6. User and Developer Experience (UDX)

6.1 Scope

The User and Developer Experience (UDX) thrust area focuses on enhancing the usability, accessibility, and productivity of PESO's tools and technologies for users and developers. This thrust area addresses key challenges in improving the interactions and workflows involving HPC software and services, making them more intuitive, efficient, and tailored to the needs of diverse user communities. By refining user interfaces, enhancing documentation, and supporting ease of deployment, the UDX effort ensures that PESO's tools are effectively adopted and deliver value to various stakeholders, including scientists, engineers, and developers. This focus aligns with PESO's broader mission of advancing scientific discovery and innovation through computational research, providing the foundation for effective tool adoption and utilization.

6.2 Team

The UDX team is composed of members from both within and outside PESO.

Lead members are:

- Lavanya Ramakrishnan: Provides external perspective and aligns cross-institutional collaboration
- Hannah Cohoon: Leads efforts in user research and engagement
- Mike Heroux: Oversees strategic initiatives for the UDX efforts

Team members include:

- Drew Paine
- Lois Curfman McInnes

2024 PESO Project Report

- Patrick O’Leary
- Wyatt Spear
- Ulrike Yang
- Tammy Dahlgren
- Xiaoye Sherry Li
- David Bernholdt

These team members contribute expertise in user research, software engineering, interface design, and user support.

6.3 Achievements

During FY24, our team focused on establishing a group of technical experts and user experience researchers to set goals and expectations for our UDX efforts. We also began work to raise awareness about user experience in the PESO and CASS communities.

- Lavanya Ramakrishnan introduced PESO members to ongoing UDX efforts at a meeting in March 2024, where she presented the STRUDEL project and its creators—the user experience team in the Scientific Data Division of LBNL.
- In June 2024, a UDX working group was established to support PESO efforts and, more broadly, CASS. This group met and identified short-term and long-term goals, including establishing a team that values and leverages user experience methods and providing resources for scientific software teams to apply user experience principles.
- Working group members (Drew Paine, Hannah Cohoon, and Lavanya Ramakrishnan) shared resources introducing user experience methods and value, which improved the working group’s understanding of UDX research methods and identified opportunities for UDX expertise to be applied.
- The July PESO project meeting breakout groups further defined the scope and goals of the UDX working group, shared user and developer experience stories and set objectives for the year.
- The UDX working group began identifying opportunities for UDX expertise by compiling examples of desired improvements for members’ projects, revealing a demand for UDX expertise on both the project level and in recognizing shared UDX challenges among projects.

6.4 Current Status

The present focus of the PESO UDX efforts is to continue learning about the challenges that scientific software projects face concerning user or developer experience. Recognizing issues allows us to plan appropriate research, interventions, and informational resources. The PESO UDX team members can participate in and leverage the efforts of other organizations, including the United States Research Software Engineer (US-RSE) association’s user experience working group. US-RSE is collecting stories about user experience efforts and challenges. By maintaining ties with the US-RSE working group, our PESO UDX team engages in the community building needed to achieve our long-term vision. Although our interests extend beyond the United States and the research software

2024 PESO Project Report

engineering community, the PESO UDX, and US-RSE UDX efforts are small enough to benefit from collaboration, particularly initially.

6.5 Plans

In FY25, we will develop a list of common UDX challenges that scientific software teams face—these may be pain points that users regularly experience or skills gaps that need addressing. These challenges will be gathered through our informal information-seeking.

We will work closely with the rest of the PESO team to improve UDX and thus demonstrate its value in software stewardship. The working group will plan a UDX research study to investigate the issue and identify possible solutions. One of our key focus products for a UDX experience study and improvement is E4S. This will benefit the studied software, and we can promote the UDX work and its impact within the scientific software community through blog posts and presentations. By raising awareness of this initial PESO UDX effort through, for example, the Better Scientific Software (BSSW) blog or one of the regular virtual US-RSE meetings, we will leverage this study to show the value of UDX investments and encourage other software teams to make such investments.

By 2028, we expect to have gathered and developed resources that communicate best practices for addressing common challenges, enabling software teams to mitigate UDX problems independently. Additionally, to encourage teams to engage in UDX work, we will provide several examples of UDX efforts within the PESO community and their impact (including FY25's study).

7. Impact Framework

7.1 Scope

The Impact Framework thrust focuses on exploring ways to evaluate scientific libraries and tools to help improve efforts' overall impact on scientific advancements and inform future funding required for stewardship. It is a challenging topic area since libraries and tools can have impacts in many and varied ways, making it difficult to determine a specific framework that is both flexible and fair while also guiding for improving impact and insight of the entire portfolio. Currently, the purpose of the focus area is to explore possible ways to measure and improve software and develop a framework that can be adapted to a wide range of products.

7.2 Team

- David Bernholdt: Lead for coordinating the impact framework working group.
- Patrick O'Leary: Contributor with expertise in software sustainability.
- Ulrike Yang: Involved in aligning programming system needs with impact assessment.
- Mike Heroux: Provides guidance on maximizing impact and stewardship funding.

2024 PESO Project Report

- Berk Geveci: Contributor focusing on the application of impact frameworks to community and product-level improvements.
- Bill Hoffman: Contributor to discussions on quality and impact.
- Terry Jones: Provides input on the productivity and efficiency of developers.
- Todd Munson: Helps identify best practices and evaluate their effectiveness.
- Keita Teranishi: Contributes to defining metrics for programming models and software categories.
- Addi Malviya Thakur: Involved in tailoring the impact framework to different project needs.
- Ross Bartlett: Contributor focusing on metrics and their implications for project assessments.
- Lois Curfman McInnes: Provides guidance on the impact on the scientific community.
- Jim Willenbring: Provides perspective on the use and communication of metrics.

7.3 Achievements

Milestones and Accomplishments:

- Formation of the Impact Framework team: The group was successfully established, bringing together a diverse set of experts to address the complex challenges of evaluating scientific software impact.
- Selection of the group's title: The team spent time ensuring the name conveyed an accurate but non-threatening description of the work, resulting in naming the 'Impact Framework Working Group.'
- Establishment of regular meetings: The group set up a bi-weekly meeting schedule starting in August 2024, providing a consistent platform for capturing ideas, discussing progress, and refining approaches to developing an impact framework.

Significant Outcomes:

- Designation as a CASS-level working group: The impact framework was made a CASS-level working group, emphasizing its importance in evaluating and guiding software across a broad scientific portfolio.
- Captured initial input on goals and metrics: The team has made progress in gathering community input regarding impact metrics, providing a foundation for further developing a flexible evaluation framework.

Metrics and Success Stories:

- Successful coordination of expertise: The working group has brought together experts from different domains, fostering collaboration to address the challenges of assessing the impact of scientific software.
- Developed early framework ideas: The team has discussed and outlined key areas for metrics, such as productivity, sustainability, and quality, establishing an initial direction for evaluating software impact in a structured way.

2024 PESO Project Report

7.4 Current Status

Progress Overview:

The impact framework working group is still in the early stages of its efforts. The team established its structure, selected a working group title, and defined a regular meeting schedule. Initial discussions have focused on gathering input from stakeholders and outlining key areas for metrics, providing a solid foundation for future work. While progress has been modest, the groundwork has been laid for effective collaboration and continued development of the impact framework.

Challenges and Obstacles:

The primary challenge at this stage is the inherent complexity of developing a flexible and fair evaluation framework for scientific software. Given the variety of software tools and their varied impacts, finding a balanced approach that meets the needs of different stakeholders is a significant task. Additionally, as the group is still in its formative phase, there are challenges in maintaining momentum and engaging the broader community effectively. Limited funding and competing priorities also pose obstacles to rapid progress.

Active Projects and Deliverables:

The key active project for the group is the development of an initial draft of the impact evaluation framework. This involves defining relevant metrics for assessing software impact, productivity, sustainability, and quality. The group is also focused on gathering ongoing community feedback to refine these metrics and ensure they are adaptable to a wide range of software products. Regular meetings are being used to capture ideas, make incremental progress, and ensure alignment across team members and stakeholders.

7.5 Plans

FY25 plans:

- **Developing a prototype Impact Framework:** The group will focus on creating a prototype version of the impact framework, complete with initial metrics for evaluating software productivity, sustainability, and quality. This prototype will serve as a basis for testing and further refinement.
- **Engaging the broader community:** The group will actively engage with the broader user and developer communities to gather feedback on the proposed framework. This will include conducting workshops or focus groups to understand the needs of different stakeholders and how the metrics can be adapted to various contexts.
- **Pilot testing of metrics:** Selected projects will be identified for pilot testing the metrics to assess their applicability and effectiveness. Lessons learned from these pilots will inform iterative improvements to the framework.
- **Collaboration with external partners:** The group will continue collaborating with external partners to align the impact framework with other ongoing evaluation

2024 PESO Project Report

initiatives, ensuring broader adoption and integration within DOE and related communities.

Expected Outcomes (FY28):

- **Established Impact Framework:** A fully developed and validated impact evaluation framework will be in place, providing a structured and reliable method for assessing the impact of DOE-supported software. This framework will be broadly adopted across different domains within the scientific software ecosystem.
- **Improved decision-making for funding and development:** The impact framework will guide funding decisions, ensuring that resources are directed towards projects with the most significant potential impact. The framework will also provide valuable insights into areas for improvement, helping projects enhance their quality and sustainability.
- **Strategic impact on scientific software:** The framework will contribute to the strategic advancement of scientific software by providing transparency and accountability regarding the effectiveness of software products. These insights will help set priorities and align efforts to maximize scientific return on investment.
- **Community-wide adoption and adaptation:** The impact framework will be recognized as a valuable tool not only within DOE but also by the broader scientific community. The flexibility and adaptability of the framework will allow it to be used across various types of software projects, contributing to a more cohesive approach to evaluating and improving scientific software.

8. PESO Community Engagement

8.1 Scope

The scope of the PESO project community engagement thrust encompasses a broad range of activities to create and foster a robust community of individuals passionate about scientific software and supporting career paths in developing and using software for scientific discovery. This work includes engaging with various stakeholders, including software vendors, PESO and CASS community members, software teams supported by SSOs, and the broader high-performance scientific computing community. This community includes national and international developers and users of scientific software across academia, industry, and professional organizations.

PESO coordinates and leads activities such as the Better Scientific Software Fellowship Program, in partnership with NSF and NNSA, and the BSSw.io web portal, in collaboration with COLABS. Additionally, PESO plays a key role in CASS community development, which involves partnerships with all SSOs within the CASS community. These activities are designed to attract individuals to the scientific software community and promote stable, long-term career paths that enable sustained contributions to scientific computing. PESO also plays a leadership role in organizing and coordinating events like the CASS Community BOF Days.

2024 PESO Project Report

8.2 Team

All PESO team members contribute to community engagement, given their crosscutting roles in partner SSOs, with additional collaboration from stakeholders within the CASS community and strategic partners, such as NSF, NNSA, and various SSOs. Key team members focused on community engagement include:

- Lois Curfman McInnes: Leader of PESO community engagement thrust, executive committee of the BSSw Fellowship Program
- Elsa Gonsiorowski: Coordinator of the BSSw Fellowship Program
- Erik Palmer: Deputy Coordinator of the BSSw Fellowship Program
- Suzanne Parete-Koon: Lead of the HPC Community Group
- Mary Ann Leung: Leadership team of PESO community engagement
- Dan Martin: Leadership team of PESO community engagement
- David Bernholdt: Executive committee of the BSSw Fellowship Program
- Mike Heroux: Executive committee of the BSSw Fellowship Program

8.3 Achievements

As a vehicle for broad community outreach, the BSSw.io website (<https://bssw.io>), a partnership of PESO and COLABS, published over 40 articles in 2024, disseminating best practices and insights on scientific software development.³ The Better Scientific Software (BSSw) Fellowship Program—which fosters and promotes practices, processes, and tools to improve developer productivity and software sustainability of scientific codes—selected the [2024 cohort](#) (six Fellows and six Honorable Mentions) and [2025 cohort](#) (five Fellows and five Honorable Mentions). BSSw Fellows are advancing scientific software practices across broad areas related to planning, development, performance, reliability, collaboration, and skills.⁴ PESO has secured partial funding for the BSSw Fellows Program for the next several years while working toward complete and sustainable sponsorship. We successfully negotiated with the US RSE (Research Software Engineering) annual conference to host the [2024 awards ceremony for the BSSw Fellows](#) at the US RSE conference in Albuquerque in October 2024, and we will announce the 2025 class at the 2025 US RSE conference.

PESO coordinated the CASS Community BOF Days in June 2024, providing a vehicle for outreach to the broader community about work across the consortium. Additionally, PESO, in collaboration with COLABS and other partners, continued soliciting new articles, events and curated content for the BSSw.io website to ensure the ongoing dissemination of valuable community-driven resources on topics related to scientific software productivity and sustainability. These achievements highlight PESO's integral role in fostering and expanding the community of scientific software developers and users. Slides are [here](#).

³ *Building and Sustaining a Community Resource for Best Practices in Scientific Software: The Story of BSSw.io*, R. Gupta et al, IEEE CiSE 2024, doi: [10.1109/MCSE.2024.3480808](https://doi.org/10.1109/MCSE.2024.3480808)

⁴ https://bssw.io/blog_posts?track=bssw-fellowship

2024 PESO Project Report

8.4 Current Status

We are currently guiding the 2024 cohort of BSSw Fellows in project completion and outreach for broad community impact, onboarding the newly selected 2025 cohort of BSSw Fellows, developing articles for the BSSw.io portal, and planning for 2025 BSSw Community BOF Days. PESO is also actively collaborating with partners in the CASS Next-generation Scientists Initiative and engaging in several working groups to expand community activities.

8.5 Plans

For 2025, PESO will continue its fellowship schedule as established in the past year, including partnership with NSF and NNSA. We will continue soliciting new articles and curated content for the BSSw.io website (via continued partnership with COLABS) to ensure a broad range of content. In collaboration with CASS community members, PESO will deliver the CASS community BOF sessions and contribute to the CASS Next-generation Scientists Initiative to address training needs across SSOs. Additionally, we aim to strengthen our partnerships and extend our outreach to include international developers and users of scientific software, promoting a national community of practice.

9. Conclusion

Computing is a foundational capability in almost all areas of scientific discovery. For problems that demand high-end computing resources, stewarding and advancing the software ecosystem established by the US Exascale Computing Project (ECP) is a fundamental responsibility, building on ECP investments.

The PESO project has significantly improved scientific software through collaborative efforts, community building, and targeted technical enhancements. By focusing on Spack and E4S, PESO has provided foundational tools that facilitate software deployment, reproducibility, and integration across diverse HPC environments.

Community engagement initiatives have strengthened the community's cohesion and improved sustainable career paths in scientific software development. The project's emphasis on user and developer experience has improved usability. At the same time, introducing impact frameworks provides a structured way to assess and enhance the impact of funded software projects.

PESO aims to continue enhancing software integration, expand support for artificial intelligence and machine learning workflows, and prioritize effective community engagement. By promoting software sustainability and improving usability, PESO is committed to ensuring that scientific software fully fulfills its critical role as a vital tool for innovation and discovery. The collaboration among PESO, CASS, and other partners supports the collective efforts needed for HPC and scientific software development, contributing to a more resilient and productive research landscape.